Platform Solutions
Online News for Developers

Issue 23, September 3, 1999

Feature Story
Each month we provide a feature article on key industry trends and developments. Authored by a member of Intel's executive staff, the feature offers insightful information for product development, planning and execution.

IDF
The Intel Developer Forum (IDF) conference is a semiannual, three-day conference targeted at business and consumer desktop, mobile, workstation, and server hardware manufacturers, and peripheral and software vendors working on advanced platform solutions. This special issue reviews just some of the 13 technology tracks comprising more than 100 in-depth sessions and 8 hands-on labs at the Fall '99 Conference August 31 through September 2.

Top Stories
Delivering in-depth reports on key platforms, products and technologies, our Top Stories provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Reader Services
If you're new to *Platform Solutions* and would like to receive this companion newsletter to the online version, please visit *Platform Solutions* on-line and go the "Subscribe Now" section to register and sign up for delivery. The online version provides direct links for quick access to the developer information and news reported in each issue, whether it's on Intel's Web site or industry Web sites.

We've created Platform Solutions to be your design news resource for the Intel® Architecture platform. If you can help us make it better, if you have a comment or a question or if you have a specific topic you would like to see covered, we'd like to hear from you. Just send an email to platform.solutions@intel.com.

If you do not want to receive this mailing in the future, please send an email to platform.solutions@intel.com with "unsubscribe" in the body of the message.

On behalf of all of us at Platform Solutions, welcome to the future of the PC platform!

Table of Contents

(Click on descriptions to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Feature

The Changing World of Commerce and Communications: Delivering the Building Blocks for the Internet Economy

Craig Barrett
President and Chief Executive Officer
Intel Corporation

The Internet economy is here today, but the work of building an infrastructure capable of supporting its accelerated growth has just begun. Intel CEO and President Craig Barrett takes a closer look at the challenges and opportunities that lie ahead.

The Internet economy is here today, but the work of building an infrastructure capable of supporting its accelerated growth has just begun. Here's a closer look at the challenges and opportunities that lie ahead as the development community works together to advance the Internet.

There's no question that the Internet has had a huge influence in changing our culture, our communications, and the way in which we do business—all within the unbelievably short timeframe of less than ten years. As recently as a decade ago, the Internet was a little-known network used by just a few scientists and Department of Defense people, and the World Wide Web didn't even exist on the communications landscape. It's hard to believe how quickly the online phenomenon has taken root, and how it's changing everything. Just as the expansion of railroads changed the face of commerce in the 1830s and the growth of interstate highways transformed transportation in the 1950s, the Internet and the Web are fundamentally altering the way people live, play and work. The principal difference, of course, is the speed of change that we're facing today.

By any measure, the growth of the Internet has been nothing short of astounding, and it's a growth that's going to accelerate even faster in the next few years. As we stand on the threshold of a new century, we're moving rapidly toward a world of a billion connected computers—and with it, a trillion connected dollars. The accelerating Internet economy is everywhere, and within a few years the distinction between homes and businesses with and without Internet access will no longer exist. That's because virtually *all* homes and businesses will be part of the rapidly evolving e-commerce world of online communications.

According to a recent Internet Commerce Market Model from IDC, more than 100 million users and 100 million computing devices now have Web access, representing a growth of more than tenfold in each category since 1995. By 2002, however, IDC forecasts that there will be some 300 million users and almost 500 million devices accessing the Web. And while in 1998 less than one percent of all gross domestic product (GDP) was generated over the Internet, that figure is expected to rise to ten percent of the GDP—approximately one *trillion* dollars—by 2002.

These trends combine to illuminate a strategic inflection point in the way business is conducted and people communicate. And as with all strategic inflection points, this particular one—advancing the Internet—presents its share of challenges and opportunities. That's because, in order for the Internet economy to truly flourish and live up to its considerable potential, we as an industry clearly need to develop a robust Internet infrastructure capable of supporting sustained growth. And we need to develop it at a faster pace than the one at which we've been moving to truly provide headroom for growth.

Building the Internet Infrastructure
First and foremost, it will take a concerted industry effort based on the adoption and proliferation of open specifications to lay the groundwork for the Internet infrastructure. That effort must be accelerated. It took five years, after all, to fully adopt the Universal Serial Bus (USB) as a computing platform standard; we have to drive open Internet specifications faster than that if we're to realize the full potential of the medium.

Simplicity and availability are also key. That's because, as ubiquitous as the PC may now seem to be, many people still don't have one because they think that computers are too hard to use. This is a barrier that must be overcome. Equally important, the bandwidth, services and content that define a truly robust Internet economy need to be provided and developed. Deployment is another significant issue: less than four percent of the network servers that will be needed to support the users and devices clamoring for online access by 2002 are now in place.

All of these requirements translate into wonderful opportunities for platform developers, as the pace of integration rises and semiconductor production migrates down to 0.18-micron geometry's. This continued technology innovation and leadership in process technology is paving the way for increasing levels of value-added innovation to be included in the building blocks that will be required to create the Internet Economy. By building blocks, I'm specifically referring to client platforms, a network infrastructure, and server platforms—all of which will be supplemented by associated solutions and services that sit on top of the platform foundation.

In the client platform arena, we must build a compelling and easy-to-use variety of computing devices that provide Internet access. We must deliver intelligent networks driven by an open, interoperable communications infrastructure, and we must develop server platforms that provide performance, scalability, value and flexibility. When you look at the effort and innovation it's going to take, the message for the developer community is clear: the growth of the Internet Economy is in our hands.

From the Business to the Home: The Internet Is Everywhere
Within the Internet economy, two truths are evident: all homes will become e-homes, and all business is moving to e-business. The infrastructure of the e-home, for example, will ultimately include multiple networked devices tying all rooms together via broadband connections. All of these devices, of course, will tie into a performance PC that will act as the central server and intelligence for the e-home network.

The e-home will also feature privacy and secure Internet access, a much broader use of online capabilities, and all kinds of new content and services. Our challenge as developers is to make the term "Internet User" obsolete by making *everyone* an Internet user, and rendering the term irrelevant. By providing easy-to-use Internet platforms, we can ensure that every home will ultimately be connected.

In the world of business, meanwhile, e-commerce is already taking off. The Internet is driving a convergence of voice and data traffic that's dramatically changing the face of modern communications—and with it, the challenges facing developers and service providers alike. As in the case of the e-home, the performance PC will be a staple of all e-business environments. From base applications to automated data components, interoperable communications to improved user interfaces, better knowledge management software to increasingly sophisticated content and services, there's all sorts of headroom for innovation. And it will take considerable innovation to meet the client and server platform demands of the evolving Internet economy.

IA-64: Fueling the Internet Economy
When it comes to creating infrastructure building blocks, Intel's IA-64 server and workstation roadmap identifies succeeding generations of processors that will provide the performance, integration and compatibility that the Internet economy will require. As our demos at the Fall '99 Intel Developer Forum illustrate, the IA-64 is up and running today in the form of the Merced processor, targeted to meet the high-end digital content creation and distribution needs of the workstation and server markets. Intel will be providing initial Merced samples to OEMs within the next few weeks, with volume production right on schedule for the first half of next year.

As an architecture that has been developed with the evolving requirements of the Internet economy in mind, the IA-64 is poised to assume a leading role as the new engine for e-business. It provides faster interactive transactions and high availability, which facilitate easy customization capabilities and ensure continuous service. Its ability to support large amounts of memory enables it to provide the responsiveness that Internet users want, and its encryption capabilities enable it to offer the security that Internet users need.

Armed with these advanced platform capabilities, the IA-64 provides the industry with the means to develop the technological building blocks that will help to create an infrastructure capable of supporting the tremendous growth of the Internet in the years to come. The truly exciting part of all this is that the growth of the infrastructure is in *our* hands. We're engaged in creating a world in which all homes are e-homes, and all business is e-business. I look forward to the challenges and opportunities that lie in store for all of us in pursuit of this goal, as we work together to advance the Internet in the years to come.

About the Author

Craig Barrett is the president and chief executive officer of Intel Corporation. He joined Intel in 1974 as a technology development manager, was named a vice president of the corporation in 1984, and was promoted to senior vice president in 1987 and executive vice president in 1990. Dr. Barrett was elected to Intel's Board of Directors in 1992 and was named chief operating officer in 1993, before becoming Intel's fourth president in 1997 and chief executive officer of the company in 1998.

Craig received his Bachelor of Science, Master of Science and Ph.D. degrees in materials science from Stanford University. After graduation, he joined the Stanford faculty and rose to the rank of associate professor before leaving the university to join Intel. Dr. Barrett is the author of more than 40 technical papers dealing with the influence of microstructure on the properties of materials science, and his textbook, "Principles of Engineering Materials," remains in use at universities today.

Top Stories

Code Clean Enables Software in Both IA-32 and IA-64 Worlds

Jim Howard
Senior Engineer
Intel ISV Performance Labs
Intel Corporation

In both Windows and UNIX* operating systems, upgraded application programming interfaces (APIs) allow source code to support both 32-bit and 64-bit versions of software. See how you can use them to offer "normal" and "extended capacity" versions of your applications.*

Intel's IA-64 architecture, together with 64-bit implementations of the UNIX, Windows and Novell* operating systems, dramatically increases the capacity and performance available to application programs. In these operating systems, application programming interfaces (APIs) have been upgraded to allow a single set of source code to build either IA-32 or IA-64 versions of applications. This allows you to offer normal and (on the IA-64 platform) extended capacity versions of your applications from the same source base.

For example, designers working with ECAD (electronic computer-aided design) applications can work directly with today's multi-gigabyte IC layouts instead of having to work separately with multiple, partitioned pieces of the layout.

“Code Clean” is the process of updating source code to take advantage of these operating systems’ multi-platform APIs, new and revised data types and updated function calls, and eliminating use of features that have been made obsolete. The code clean process has several elements:

- Selecting an appropriate porting model (naming, directories, data type use, and conditional compilation macros).
- Assuring regression tests perform nearly complete code coverage.
- Revising source code so it builds on both IA-32 and IA-64, including corrections for data type agreement and algorithm updates.
- Compiling updated source with both 32-bit and 64-bit compilers.
- Running regression tests in both worlds.

One value of doing code clean that is often overlooked is its ability to add reliability, robustness, and maintainability to applications on all platforms. Maintenance and support is expensive. Code clean can save you money right now.

Enabling Software in Both Worlds
The updated operating system APIs have preserved the names and semantics of many of the familiar data types, including `int`, `DWORD`, and `REAL*8` and have revised the definition of others to scale with the capacity of the target architecture. These include `long` (in UNIX), and `UINT_PTR` (in Windows). In addition, both Windows and UNIX now offer sized types, allowing exact specification of the precision of any data value in your source:

In Windows: `__int32`, `unsigned __int16`, ...
In UNIX: `int16_t`, `uint64_t`, ...

This combination of types and `typedefs`, known as the platform's Data Model, allows revision of data types used in source code to offer the desired features and capacity on either platform.

Certain features have also been deprecated, or scheduled for obsolescence. These features are being made obsolete because they could not represent the larger capacities of IA-64, or could not accept or supply 64-bit sized quantities on 64-bit platforms.

For example, the Windows API call `GetClassLong()` returns a 32-bit sized value that could be a pointer or a number, will still work on 32-bit Windows platforms and is not available on 64-bit Windows platforms. A replacement API call, `GetClassLongPtr()`, performs the same function and returns a platform-bit-size scalable (polymorphic) value. The new API works like before on 32-bit platforms and returns 64-bit sized data in Win64 on IA-64. Revising source code to use only the new API allows it to run properly on both platforms.

Windows and UNIX Make Different Choices
The UNIX and Windows communities each chose API and data model revisions they believed would minimally impact their existing users. While both communities adhere to ANSI C, the commonly used `typedefs` of each operating system led to different choices for the types that scale with architectural bit size. For example, the UNIX vendor community chose at the “Aspen” conference in 1996 to make `long` scale with the architecture's pointer size. Microsoft chose to preserve `long` as 32-bits, and make other `typedefs` scale, including `INT_PTR`, `LONG_PTR`, `DWORD_PTR`, and so on.

If source code must be built for both IA-32 and IA-64, you must avoid use of the `long` data type. You can do this in either of two ways:

By using ANSI types that scale properly (like `intptr_t`).
Or by using an invented type like `int3264_t` that is a conditionally compiled `typedef` in a "compatibility" header file that gets `#included` everywhere.

Code Clean of Data Type Use
The following API revisions comprise the majority of code clean:

- Updating code that mixes pointers and integers: change `(int)`*`pointer`* casts to `(uintptr_t)`*`pointer`*.
- Fine-tuning variable and structure field sizes: don't waste space using an `__int64` or `long` if an `int` will do.
- Re-optimizing field ordering in structures: minimize padding (wasted space) by placing scalable types early in the `struct` definition.

Assuring Software Correctness
Correcting type scalability (size polymorphism) is only part of the code clean process; it isn’t possible to automate it all. Only a human can decide when algorithms must be updated and avoid undesirable increases in data structure sizes.

Avoiding "Data Bloat"
Pointers in the 64-bit world are double the size of those in the 32-bit world. If key, often-used application data structures contain mostly pointers, such records will double in size, increasing the load the application presents to the system. Finding alternate, smaller ways to express pointer-like relationships among records can dramatically reduce “data bloat” and improve overall application performance.

Some possible mechanisms include:

- Making arrays of records instead of pointer chains.
- Using base-plus-displacement addressing (32-bit displacements).
- Using 32-bit `__based()` pointers (in Windows only).

Casual use of `int` and `long` is very common because both have been 32-bits and are of similar cost. This is no longer true. Application source should be analyzed for appropriate data type use:

- Use the smallest type necessary for the datum or field.
- Check function parameter types.

Algorithm Updates
Many algorithms need "tweaking" in the presence of 64-bit platforms. These updates can include:

- Smarter heap allocation routines which place records of pointer chains close together.
- Data structure walking to take advantage of the above.
- Update of algorithms that hash 64-bit pointers or integers to adequately "randomize" those bits actually used.
- Fixing any mixing of `double` and UNIX `long` or Windows `__int64`, which can lose bits of precision.
- Differences in shifting and masking using hex constants.

Build Rule Updates
The file system directory trees used to hold source code and generated (i.e., compiled) files, plus build rules themselves, will need to be updated to represent multiple platforms. This can be accomplished in several ways:

- Defining platform-named directory trees to receive each platform's build results.
- Making files suffixes parametric; `$(OBJ)` is "`.o`" in UNIX, and "`.obj`" for Windows.
- Using `makefile includes` or sub-projects to pick up platform-specific definitions.
- Assuring compilers, etc., are invoked with appropriate "`/D`"(efines) and target locations.

You Should Begin "Code Clean" Now
The market demand for high-capacity versions of many applications emphasizes the value of beginning code clean *now*. While this article is necessarily abbreviated, other, more detailed resources are available to aid doing code clean, including:

- Attending the IA-64 track at the Intel Developer Forum.
- The Microsoft Windows Platform SDK, which contains a 64-bit C/C++ type checker.
- Project Monterey's 64-bit "lint" to check out your code.
- Code coverage analysis tools from many sources to gauge test suite adequacy.
- Check out the clean code Web-based training that will be available on Intel's IA-64 homepage in September.

A good approach is to start with a favorite tool (whose source is available to you) and code clean it today.

About the Author
Jim Howard is a senior engineer in Intel's ISV Performance Labs, which is devoted to helping software designers take best advantage of Intel technology for performance or capacity. Jim is currently working on methodologies for migrating applications to IA-64. He was previously engaged in the development of ECAD tools.

For More Information

- See Intel's IA-64 Web page.
- Online Web-based tutorials are available at the Intel Web site for the IA-64 instruction set, and the Code Clean process.
- The Windows 2000 SDK is available from Microsoft.
- UNIX LP64 programming information is available.
- IA-64 programming information for most every OS is available. The Intel web page provides further reference links for Windows, multiple UNIX implementations, and Novell.

The Intel® 820 Chipset and the Next-Generation Performance Platform

Frank Hady
Senior Staff Systems Architect
Intel Architecture Labs
Intel Corporation

The new Intel® 820 chipset delivers the concurrent bandwidth and low loaded latency needed to keep pace with today's faster CPUs and graphics accelerators. But what can it do for application runtime? Recent test results from Intel Architecture Labs provide the answers.

The new Intel 820 chipset delivers the concurrent bandwidth and low-loaded latency needed to keep pace with faster CPUs and graphics accelerators. But what can it do for application runtime? Recent test results from Intel® Architecture Labs provide the answers.

Faster CPUs and graphics controllers in high performance desktop PCs place higher bandwidth demands on the chipset. To keep pace, the chipset must feature higher-bandwidth ports. The Intel 820 chipset meets this requirement with the 133 MHz front-side bus (FSB) and AGP 4X, providing a 1 Gbyte/sec. port for processor connection and a 1 Gbyte/sec. port for graphics.

But how does the 820 chipset actually affect application runtime? To show how, Intel Architecture Labs conducted a series of tests, beginning with evaluations of concurrent bandwidth delivered by several chipsets, including the Intel® 440BX AGPset and the 820 chipset.

As Figure 1 shows, the ability to keep up with the bandwidth demands of high-performance processors and graphics controllers requires the additional main memory bandwidth available from RDRAM, as well as a chipset that can deliver low loaded latency performance under heavy application demands.

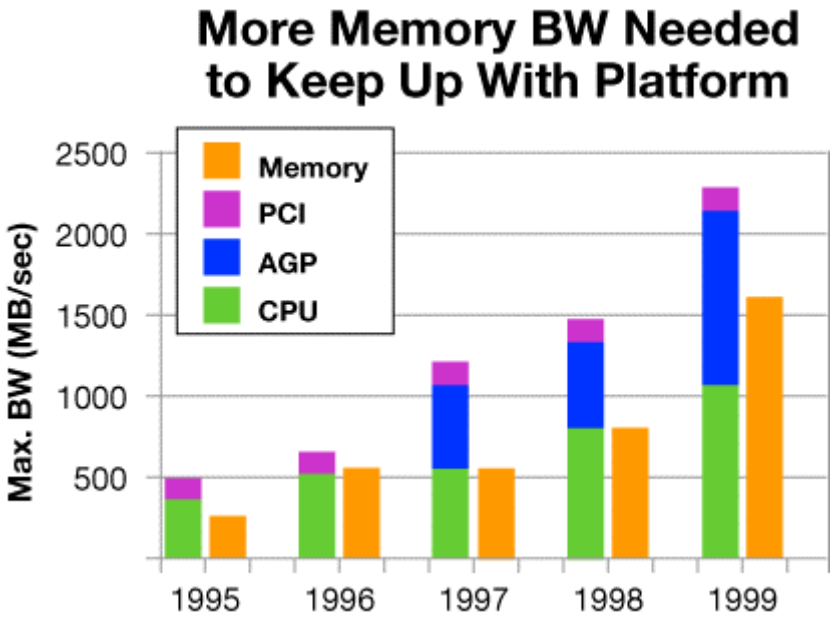

Fig.1 More Memory Bandwidth Needed to Keep Up with Platform

Figure 2 shows the concurrency envelope for four different chipsets. Lines are traced in the lab by varying the FSB and AGP loads, with the resulting curve showing the region of concurrent CPU and AGP bandwidth in which the application must operate. The 820 chipset provides are region almost four times the size of the 440BX AGPset region. Before the 820 chipset, the 440BX was the concurrency champ, providing the largest concurrent envelope.

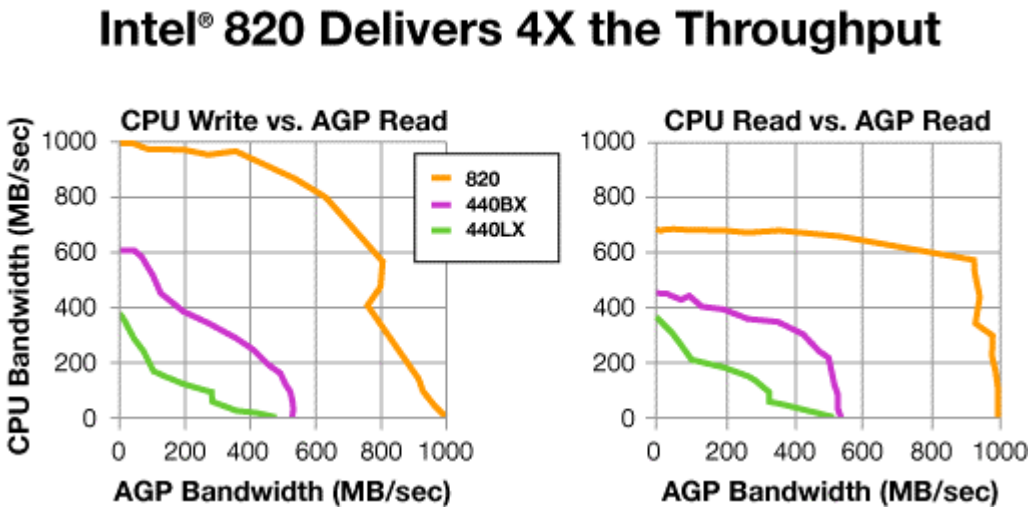

Figure 2. Intel® 820 Chipset Delivers 4X the Throughput

New Applications Push the Bandwidth Envelope

Yesterday's applications, including word processors and office suites, do not push the envelope of current chipsets. Their application loads feature low FSB utilizations (often less than 10%), typically with just one outstanding CPU request, and little contention from AGP or PCI. With low FSB utilization, there is little the chipset can do to enhance application performance.

Benchmarks that test one-year-old 3D graphics applications tell a radically different story. These benchmarks show bandwidth demands that push beyond the concurrency boundary of the 440BX AGPset. Figure 3 shows measurements using "Toonenstein,"* a soon-to-be-released 3D application. The Toonenstein bandwidth histogram is overlayed with the chipset concurrency charts. The histogram has the same axis as the chipset concurrency charts and shows the distribution of 300 separate samples of concurrent chipset bandwidth (50usecs each) measured while Toonenstien was running and quantized into buckets of 25MB/sec in size along both axis. Green regions show combinations for AGP and CPU bandwidth found in no samples, a yellow region shows a bucket where 0 to 0.5% of the samples fall, dark yellow shows .5% to 1%. The sum of all the buckets in the chart is 100%. This application greatly exceeds the 440BX envelope, spending 20 percent of its time outside the 440BX concurrency envelope.

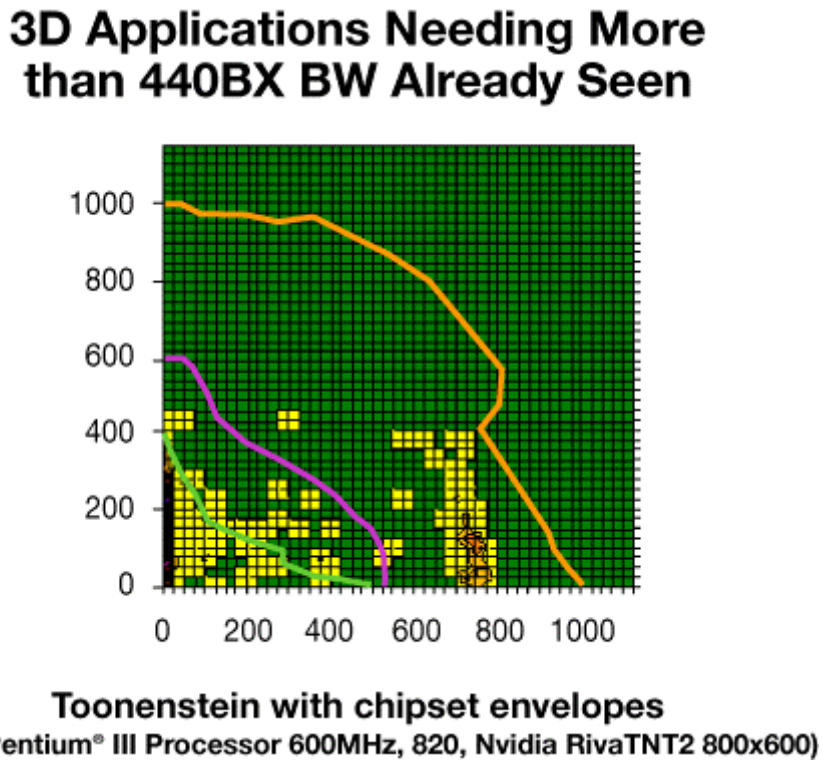

Figure 3. 3D Applications Needing More than 440BX Bandwidth Already Seen

Why Sample Duration Is Critical

IAL test results also show why relying on average values alone for your chipset bandwidth analysis can be quite deceptive. This is because applications typically exercise very high concurrent bandwidth over very short time periods. The highest bandwidth usage can be concealed when tests solely rely on average bandwidth values.

Figure 4 shows that the bandwidth characteristics of applications can appear to vary greatly as sample times are changed (all previous charts used 50 µsec samples).

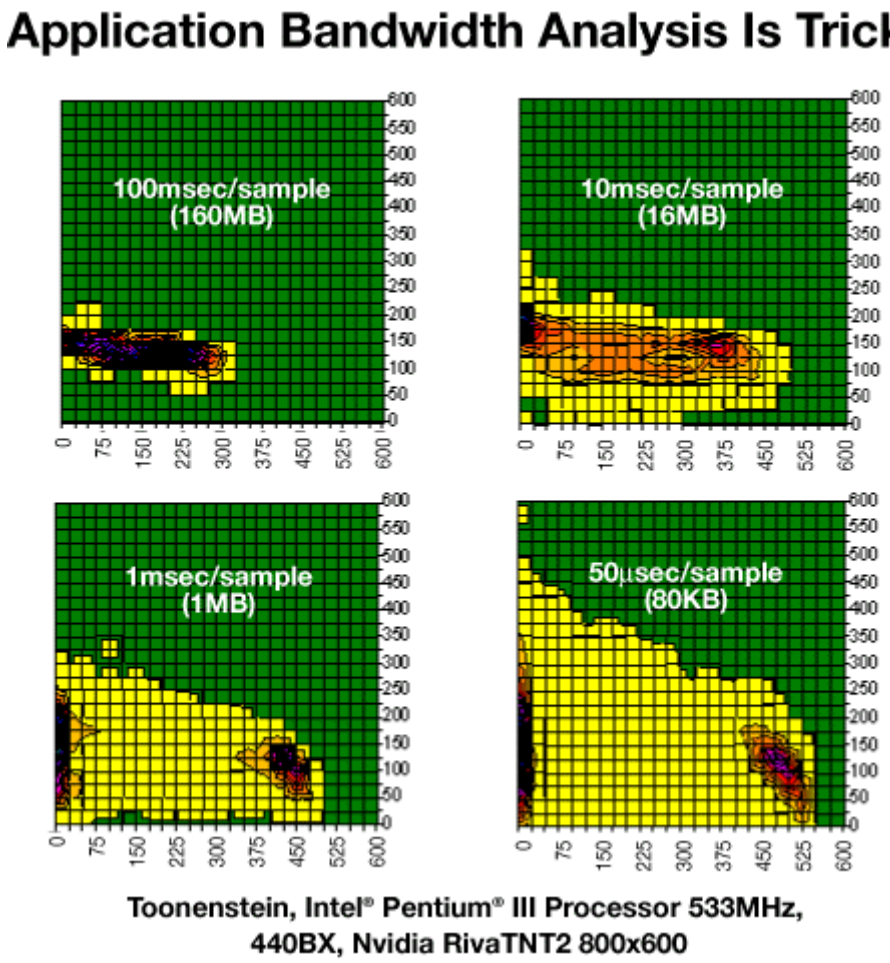

Figure 4. Bandwidth Analysis Is Tricky

What is the appropriate sample duration? The answer depends on the characteristics of the application being measured. Since textures are big, large sample sizes (50µsec samples) can work well. As Figure 4 shows, to see what's really happening with processor bus contention, you need to look at much smaller sample sizes.

Read Latency Affects Application Runtime

When we look at applications that make infrequent use of memory, the chipset can do little to effect runtime. By comparison, the chipset plays an important role in determining the runtime of more memory-hungry applications. In most cases, the role of the chipset is to return data (32-byte cache lines) rapidly to the processor, minimizing the time the processor spends waiting and not working.

In high-traffic applications, the "loaded read latency," which is defined as the main memory read time when the chipset is busy has a measurable impact on application runtime. In addition to featuring a high level of FSB traffic, such applications may also involve a high volume of AGP activity.

The 820 chipset is designed to reduce loaded latency and avoid wasted CPU clock cycles four ways:

- Hiding concurrent AGP traffic.
- Hiding flushing of internal write buffers.
- Hiding DRAM refresh.
- Minimizing DRAM page misses.

Test results with the StreamNT Triad* test, as shown in Figure 5 confirm that the 820 chipset delivers significantly lower loaded latencies than the 440BX AGPset, resulting in a faster runtime for the application.

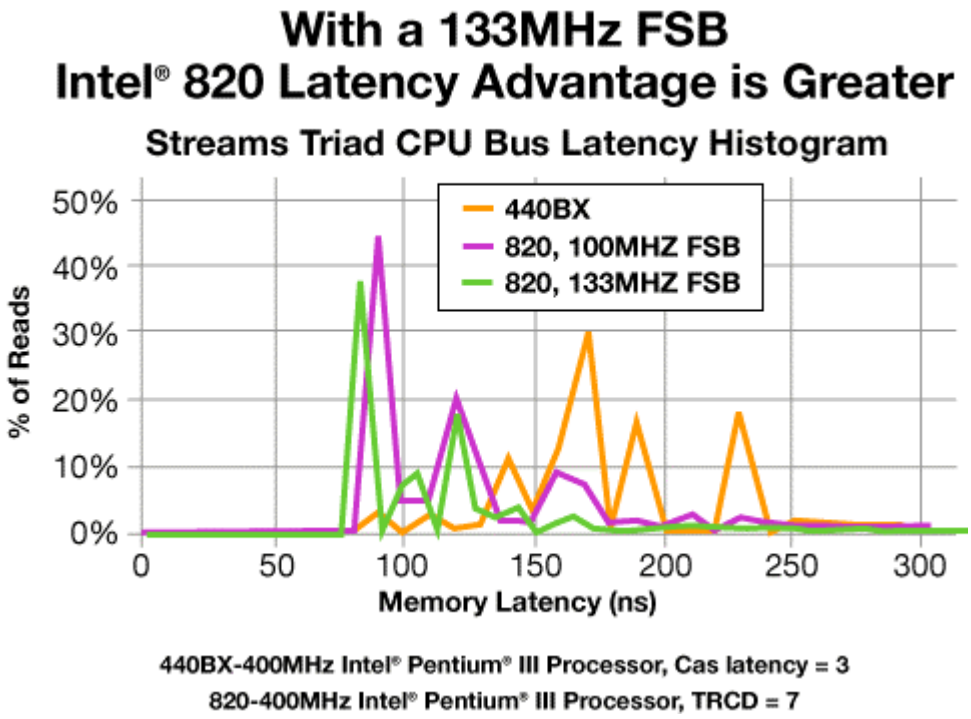

Figure 5. Intel 820® Chipset Proves Lower Loaded Latency

Get Ready to Meet Higher Bandwidth Requirements
Faster processors and graphics accelerators have a tendency to enable better applications, that in turn drive greater bandwidth requirements. This trend can be expected to continue as developers seek to add even greater realism to their 3D software titles.

For example, the use of smaller triangles, as measured by the 3DMARK* small triangle (6-pixel) test, enables developers to add more detail, realism, and complexity to scenes in 3D applications. Compared to the 440BX AGPset, the 820 chipset shows 27 percent faster performance on this test.

Test results show that designing PCs with the 820 chipset can make a real difference in runtime for high-bandwidth applications. This performance gain provides developers with the flexible high-performance platform they need to create more lifelike and compelling 3D performance in the next generation of software.

About the Author

Frank Hady is a senior staff system architect in the I/O Group at Intel's Platform Architecture Labs, where he works to improve the I/O capabilities of the PC. Prior to joining Intel in 1995, Frank researched parallel computer network performance at the Supercomputing Research Center. He received BSEE and MSEE degrees from the University of Virginia and a Ph.D. in electrical engineering from the University of Maryland.

For More Information

- Intel® 820 Chipset
- Intel® 440BX AGPset
- Streams NT Triad

USB 2.0: An "Evolution" Underway

Jason Ziller
Technology Initiatives Manager
Desktop Products Group
Intel Corporation

USB is growing in popularity, but one limitation is speed. Maximum throughput is just 12 Mbits/sec. To speed things up, Intel and other members of the USB Implementers Forum are launching USB 2.0, a next-generation specification that boosts USB's throughput by a factor of 20.

USB (Universal Serial Bus) has gained powerful momentum among PC OEMs, peripheral vendors, and users. Because USB is easy to use and compatible with a wide range of peripherals, virtually all new PCs feature USB ports, and most peripheral vendors now include USB in their new products. Now, Intel and six other industry-leading companies (Compaq, HP, Lucent, Microsoft, NEC, and Philips) have launched USB 2.0, a next generation specification that will boost USB's current maximum data throughput by a factor of 30 to 40.

Earlier in the year, Intel and the USB 2.0 Promoter group announced a speed target of 10 to 20 times faster than USB 1.1. However, at the Intel Developer Forum Fall '99, the group revised the target up to 30 to 40 times faster than USB 1.1. The decision resulted from engineering studies that concluded 360 to 480 Mbs could be achieved without adding cost and complexity over the original target. This speed range is expected to be sufficient to satisfy the bandwidth requirements for the most demanding user applications.

USB 2.0 will extend USB's benefits to peripherals such as video conferencing cameras and high-speed printers, and will allow users to connect multiple high-speed devices simultaneously. The USB 2.0 specification draft is planned for release at the USB 2.0 Developer Conference in October. USB 2.0 systems and peripherals are expected to be shipping by the second half of 2000.

Ease of Use
USB 1.1 has gained popularity among OEMs and independent hardware vendors (IHVs) because of its low cost, plug-and-play capability (including dynamic attach and detach), and easy expandability. USB is also a key element of Intel's and Microsoft's Easy PC initiative designed to make PCs easier to set up and use. USB 2.0 will extend this ease of use to even more peripherals by providing a low-cost, easily deployable bus that can connect PCs to multiple high-performance devices and enable several of them to run concurrently.

Transition to USB 2.0
It's important to understand that USB 2.0 is an evolutionary advance that retains all the existing benefits of USB 1.1, including plug-and-play, dynamic attach and detach, power distribution to devices, and support for power management. USB 2.0 provides backward compatibility with USB 1.1 and will be transparent to end users. In fact, all existing USB 1.1 devices will have full compatibility with a USB 2.0 system. In addition, today's USB cables and connectors are expected to support USB 2.0 without modification.

The end user benefit of USB 2.0 is significant. It will enable the use of higher performance and higher functionality peripherals such as high-resolution video cameras, high-speed printers and scanners, fast storage, and broadband Internet connections. USB 2.0 has enough bandwidth to support the most demanding user applications, such as image editing, video editing and interactive gaming, where multiple high-speed peripherals will be running simultaneously. USB 2.0 can also make today's applications much more productive. For example, the time it takes to download 40 MB of photos will be reduced from about 6 minutes on USB 1.1 to only a few seconds on USB 2.0.

Implications for Developers
USB 2.0 will provide an inexpensive way to connect PCs to high-performance peripherals, and PC OEMs will be able to add this capability to their systems with little impact to total system cost. The ubiquitous USB 1.1 connectors present on today's PCs are expected to become USB 2.0 connectors in the future.

The news is equally positive for peripheral manufacturers. USB 2.0 will expand the market segment for USB peripherals. In terms of engineering effort, designing a USB 2.0 peripheral will be similar to designing to the USB 1.1 specification. This is a compelling opportunity for peripheral vendors, since these same peripherals can be made to connect to the huge installed base of USB-capable PCs as well.

Developers should prepare now to implement USB 2.0 in the year ahead. Plan to attend the USB 2.0 Developer Conference in October, where the USB 2.0 specification draft will be released. A USB 2.0 Adopters Agreement is required to attend the conference and receive the specification draft (see below for the Web site that contains the agreement). The final specification is scheduled for release in the first quarter of 2000, just a few months from now. The specification is expected to lead to the development of USB 2.0-compatible PCs and peripherals, with products available in the marketplace starting in the second half of 2000.

For More Information

- Visit the USB Implementers Forum Web site for information about the availability of the USB 2.0 specification, including the required adopter agreement.

- Download the new white paper, An Overview of USB 2.0 from the USB Implementers Forum Web site. It recaps USB 1.1 and introduces the features and benefits of USB 2.0.

About the Author
Jason Ziller has held a variety of technical and marketing positions at Intel over the last 14 years. The first half of his Intel career was spent in engineering on flash memories and microprocessors. Later, Jason held technical marketing, strategic marketing, and platform marketing management roles in the Mobile Microprocessor Group. Jason is currently technology initiatives manager in the Desktop Products Group in charge of USB. He also chairs the USB Implementers Forum.

Introducing the Intel® Internet Exchange (IX) Architecture

Mark Christensen
Vice President and General Manager
Network Communications Group
Intel Corporation

A truly useful network platform should offer a range of performance, the flexibility to handle multiple protocols, and world-class development tools. The Intel® IX Architecture is the first industry design methodology to meet all three goals.

Rapid evolution of network systems driven by the ubiquity of the Internet economy requires a new vision of open, versatile, programmable silicon and software solutions. The Intel® Internet Exchange Architecture (IXA) defines a framework for network system OEMs to meet the Internet economy challenge. It offers them a comprehensive set of building blocks to define a broad range of platform solutions.

The IXA is among the first industry architectures to offer an open development environment, one that helps network system OEMs to select interoperable, system class solutions from Intel third-party developers and internally developed solutions.

Why a New Architecture Now?
The impetus for creating IXA comes from the inflection point at which the networking industry finds itself today. Only three or four years ago, it was a rule of thumb that 80 percent of the traffic on a network was local, while 20 percent was outside the four walls of the enterprise. Thanks to Internet phenomena such as the widespread use of email and web-based applications, that pattern is now reversed, with 80 percent of traffic taking place outside the enterprise.

This means packets are traveling farther and across more complex infrastructure, so processing them has become much more complicated. They must be routed more efficiently and at various quality-of-service levels. As a result, network switches—which only a few years ago were focused on optimizing layer 2 functions—now must enhance more sophisticated layer 3 and layer 4 functionality as well. Of necessity, the network is becoming more intelligent, with decision-making capabilities migrating from upper levels down toward the wire.

For the network equipment developer, it is no longer a question of building a LAN switch or a WAN switch: the line is disappearing between LAN and WAN. Today's switches must deal with many protocols, provision and meter many classes of service, and span the enterprise, access, carrier, and ISP networks.

IXA and its centerpiece, the Intel® IXP1200 network processor, are designed to meet these diverse needs. We have assembled a wide range of solutions that developers can program to meet their requirements, including a variety of third party solutions, and we have designed these elements to work together.

Elements of the IX Architecture
The IXA addresses a broad range of Internet system infrastructure with interoperable solutions. The Intel® Pentium® III processor is a leading web applications processor of layers 5-7. IXA fills in the building blocks at layers 1-4 of the OSI model.

IXA is designed to give the network equipment developer a cohesive set of hardware and software building blocks:

- Intel® IXP Network Processors—combines six onboard Microengines and an ARM® architecture compatible core processor. The Microengines handle computing that otherwise would have to be done in ASIC, providing wire speed processing in a much more programmable, flexible model. The StrongARM® core processor is typically used to run the control layer for the networking box.

- Intel® IXE Application Engines—these building blocks can be used to provide high performance switching functionality for ATM, Gigabit, and 10/100Mbps Ethernet.

- Intel® IXF Formatting Devices—a broad range of devices is available to format ATM cells, T1/E1 frames, Sonet/SDH frames, and Gigabit and 10/100Mbps Ethernet packets.

- Intel® LXT Physical Interfaces—these include interfaces for T1/E1 lines, HDSL, HDSL2, Sonet/SDH PHYs, 10/100Mbps Ethernet, and Gigabit Ethernet (including Gigabit over copper).

- Intel® IX Platform Development Environment—Intel offers a test-bench environment that allows the developer to do the programming for a solution well before hardware development is completed.

Toward the Intelligent Network

The biggest challenge today is building adaptable intelligence into the network at all levels without sacrificing throughput. For example, today's IP packets, with their embedded classes and qualities of service, often require performing 128-bit header inspections while maintaining line speeds in the OC-12 range or higher. IXA is designed to answer this challenge by providing programmability and adaptability below the application layer.

Our IXP1200 network processor performs primarily in layers 3 and 4, specifically for switching and packet prioritization. The IXE application engines provide additional transport-specific layer 3 functionality. The IXP1200 is also recommended for bridging between different types of transports—for example, at a LAN/WAN edge, taking in an Ethernet packet, routing, constructing an ATM cell, and sending it out to the Internet.

When the network processor sends out a cell, it travels over our new high-speed IX bus. The IXP1200 offers both a PCI bus for connecting with control and management functions, and the IX bus for attaching layer 2 IXF devices with minimum overhead.

A New Definition of "Platform"

The IXA is more than a multi-layer collection of silicon components. We believe it represents a new way to define a product development platform for networking systems.

Today's networking systems can be much, much broader. A narrow, top-down or bottom-up focus is no longer sufficient. OEMs are moving from thinking in terms of product specifications—for an ATM switch, for example, to thinking in terms of an architecture encompassing multiple protocols.

It's now time to take the next step, from architecture to platform. True platforms are open, standards based and highly versatile. They also give the developer four key capabilities: scalability, product differentiation, faster time-to-market and lower total cost of ownership. How does IXA measure up to this platform definition?

- Scalability—using our Intel® IXB3208 IX Bus Scaling Fabric, network system OEMs can connect multiple IXP1200s to deliver up to 10 Gbps of system performance. With IXA, developers have an architecture that can be just as applicable to an ATM port aggregator as it is to a VPN blade. This is a major point of differentiation when comparing IXA to other development platforms.

- Reduced cost of ownership—scalability also reduces total cost of ownership, because it enables an OEM to extend the life of a product while preserving their investment in architectural development. For example, a product can be moved from a single-chip first generation to subsequent "multiprocessor" generations as the market matures. Another factor reducing cost of ownership is that the OEM is able to start with a programmable building block, avoiding the high costs of custom ASIC development.

- Product differentiation—with six programmable Microengines and a programmable core processor, OEMs have plenty of opportunity to create the specialized algorithms and proprietary programs that differentiate them from the competition. The OEM owns that differentiation no less than if it were hard-wired in a custom ASIC.

- Faster time-to-market—building a custom ASIC typically involves a 12- to 18-month development cycle. If protocol standards change or if additional services become necessary—certainly a frequent occurrence in the networking arena—the OEM faces a very painful updating process. IXA building blocks help OEMs to bring better products to market, supporting smarter bandwidth and multiple protocols within much shorter development cycles.

Conclusion

According to feedback Intel has received from network system and subsystem developers, the IXP1200 network processor appears to be the most versatile and cost effective available on the market today. The experience of beta OEMs has confirmed a significant price/performance advantage. But perhaps the most unique and valuable attribute of our network processor is that it's part of the complete Intel Internet Exchange Architecture.

IXA offers network system builders a framework to define succeeding generations of network equipment. OEMs can select silicon building blocks for a range of enterprise, access, and carrier class networks. They can select a whole new design methodology based on the IXP1200. And they can utilize complete platforms based on hardware and software reference designs to deliver the ultimate in time-to-revenue.

About the Author

Mark Christensen is Vice President and General Manager of Intel's Network Communications Group (NCG) headquartered in Hillsboro, Ore. In his current position, Christensen is responsible for the management of Intel's networking business, including operations located in Hillsboro, Salt Lake City, Utah, Hudson, Mass., Copenhagen, Denmark, Haifa, Israel, and Jerusalem, Israel.

Christensen joined Intel in 1982 as a Manufacturing Development Engineer for Oregon Sites System Manufacturing. Prior to his current position, Christensen served as Project Manager for Intel's network adapter line introduction, and as the Director of Marketing for the Network Products Division. Mark was promoted to Vice President and General Manager of Network Products Division in 1997. Since taking over as General Manager, Christensen has sponsored multiple growth initiatives to expand Intel's networking focus including the acquisitions of Case Technologies, Dayna Communications, Digital Semiconductor (Network Products Division), and recently Shiva Corporation. He was promoted to Corporate Officer in January 1999.

Christensen was born in Fremont, Ohio, in 1959. He received a Bachelor of Science degree in Industrial Engineering from Oregon State University in 1982, and a Master's degree in Business Administration from the University of Oregon in 1989.

For More Information

The Intel Internet Exchange Architecture represents a new definition of the development platform, and a solid foundation for building the network systems of the Internet economy. For more information including application profiles, please visit the Intel Networking site.

Note: ARM and StrongARM are trademarks of Advanced RISC Machines, Ltd. Third-party brands and names are the property of their respective owners.

Preparing for 2001 with a New Power Supply Architecture

Bill Colson
Marketing Manager
Intel Architecture Labs
Intel Corporation

New high-performance desktop processors from Intel are on the industry's radar screen. By 2001, the ATX motherboard standard will no longer be able to meet the power supply, EMI and thermal requirements of systems using these powerful chips. Now is the time to prepare new solutions.

A new generation of high-performance desktop processors from Intel is on the industry's radar screen. By 2001, the venerable ATX standard will no longer be able to meet the power supply, EMI and thermal requirements of systems based on these powerful new components. Now is the time to prepare new solutions. This article introduces one key building block, the Intel® New Power Supply Architecture.

It seems hard to believe, but the current ATX form factor has already been an industry standard for five years. While it will continue to be a valid desktop PC design solution, this will change with the advent of new high-performance Intel processors slated for launch in 2001. Today's power supply, EMI and thermal solutions will not be adequate to support these next generation chips, which include both IA-32 and IA-64 processors.

Improvements to the ATX infrastructure will be necessary for the developers to support these new chips while meeting performance and cost targets for high-capacity desktop computers capable of delivering new broadband, speech, and vision technologies. For want of a better term, let's call this solution the "next generation" chassis, running on a New Power Supply Architecture (NPSA). This is the first of a series of articles that explore the elements of this technology. This article provides an introduction to the benefits of NPSA.

The Fall '99 Intel Developer Forum includes course material that supports the Scalable Platforms Initiative. Attending IDF is the best way to gain technical information on the platform technologies needed to meet near-term challenges on the desktop. The "Designing Next Generation Performance Desktop Solutions" track will provide you detailed information on NPSA and how it reduces platform cost and physical size while increasing efficiency.

Pushing the Power Supply Envelope

Power supply architecture hasn't changed in decades. While traditional PC power delivery is not highly efficient, it has not cost much and has provided adequate performance. Incremental improvements in power supplies have occurred to balance cost with performance, and these power supplies are still acceptable for today's PCs.

With the advent of new processors, chipsets and graphics controllers, time is rapidly running out for legacy power supply architecture. Remember that just two years ago system power loads were in the 135-watt range. Today, system power requirements are closer to 250 watts, pushing the envelope of current power supply technology.

Next generation systems will require 500 watts. Peak current (Icc) and maximum current slew rate (dI/dt) are increasing, while processor supply voltage tolerances are decreasing. It is abundantly clear that present power supply technology cannot handle these requirements. NPSA is an idea whose time has come.

NPSA: Simpler, Cheaper, and More Efficient
The logical conclusion is that since we need to change power supply technology anyway, it makes sense to design a new power supply architecture with fewer parts that eliminates redundant AC/DC conversions. NPSA is conceptually simpler, easier to upgrade, and inherently more scalable to meet the requirements of next generation platforms.

NPSA offers clear advantages over legacy power supply technology. Intel Architecture Labs is working to achieve the following design goals:

An overall cost reduction in the area of 30 percent—50 percent cost savings compared to the existing system power supply box and 20 percent savings over present voltage regulator (VRM) design.

- Up to 80 percent efficiency—compared to 65 percent efficiency for conventional PS/2 power supplies, while maintaining the same form factor and size of existing units.
- Improved reliability with fewer parts—up to 20 percent improvement in usage of form factor area.
- Power Factor Correction (PFC) included at no extra cost.
- Cooling by the integrated 120mm-system fan—with no internal fan required in the power supply unit.
- Faster transient response—almost instantaneous voltage rise-time without “ringing.”

NPSA Status Update
Intel Architecture Labs (IAL) has completed NPSA proof of concept implementations at 0.5 MHz, 3.3 MHz and 5 MHz in several topologies and demonstrated a PC with an Intel® Pentium® III processor powered by 3.3 MHz NPSA-based system.

The NPSA 0.5 specification is currently slated for release in the first half of 2000.

Together with other improvements, NPSA is a key enabling technology that will help the industry meet customer expectations for advanced high-performance desktop systems in 2001. Now is the time to learn more about this “high-powered” lower-cost alternative to legacy power supplies. IDF is the place to start.

About the Author
Bill Colson is a 15-year Intel veteran. In his role as a marketing manager in Intel Architecture Labs, he is primarily responsible for platform architecture technologies. He holds a patent in the area of server management and is a member of the IEEE and Next Generation I/O Forum. Bill has written articles for EE Times and Electronic Design. He has presented at numerous industry forums and tradeshow events. Bill holds a B.S. in computer systems engineering and electrical engineering from the Oregon Institute of Technology.

For More Information
Attend the Fall ’99 IDF for technical information on NPSA. Check here for regular updates on NPSA, NuATX and other key building blocks of Intel’s Scalable Platforms Initiative.

Video As Input: An Opportunity Knocking

Video as Input (VAI) technology treats the video camera as an intelligent input device for the PC, rather than a simple streaming source. This input can allow the PC to sense and react to environmental changes. As a result, new user experiences could be coming into view as early as the year 2000.

Video As Input (VAI) is an entirely new category of digital entertainment capability that delivers new computing techniques by addressing every industry level, from base platforms to end-user awareness. The Intel Architecture Lab (IAL) is working with associates that include PC game creators, toy manufacturers, and camera vendors to enable a class of applications that will support the PC + Camera to deliver new value to the end-user. At this juncture the full potential of VAI has not been tapped and a synergistic effort to develop the VAI capabilities will enhance the future of consumer VAI.

As part of the development of VAI capabilities, the Digital Entertainment (DE) initiative within IAL is developing segmentation, tracking and recognition technologies that can enable several new forms of interactive, digital play. Let's discuss how this emerging technology is being used today and the impact it can have in the future.

Introduction to VAI
The traditional trends of computing are output focused. Input devices, however, have barely been investigated. The original keyboard that was developed many decades ago is still the primary means of input today. With the appearance of the mouse, joystick and touch screens, the computer increased very little in awareness of its physical surroundings. In the last few years, however, microphones and video cameras have begun to be shipped with new computers, enabling a fundamental change in how computers can discern their local environment. Because humans view the world primarily by hearing and seeing, a shift away from more "touch" based input devices is needed. The PC does today, however, rely almost entirely on touch-like input devices.

Currently, most video technologies for the PC reside in codecs, conferencing, television, and media display. The amount of intelligent, semantic-based processing applied to the video stream is insignificant. The PC takes video in, changes representation (color conversion, compression), and displays it on the screen.

Video as input is about exploring the use of higher levels of cognitive examination of the video stream. It is about treating the camera (the eye) plus the PC (the brain) as an intelligent input device, rather than a simple streaming source. This input can allow the PC to react to environmental changes.

Some examples of how VAI applications can be used are:

1. Proximal Sensing/passive input—passive monitoring with sophisticated event triggering based on recognition, this technology could be used in security, home-office automation, etc.

2. Augmented Reality—capture video input, augment it and then display it back to the user, this could be used in games or immersive programming where the user inserts self into virtual locales (such as a ski slope or music video) and interacts with objects in virtual space.

3. Human-Computer Interface—integrates video into the primary interface to the PC and appropriate applications. A set of visual UI control primitives like buttons and sliders in Windows could be used as touch screens with standard monitors, "graspable" UI, and gesture-driven command and control for selected applications.

4. "Virtual Reality" based on video input—synthesize output (3D models and photo-realistic video) based on video (and other semantic) input – where emotions detected in speech would synthesize "avatars" or drive 3D models of virtual worlds or gesture-driven command and control gaming applications.

Desktop cameras are becoming ubiquitous, thanks to video conferencing and imaging applications. It is now easy to add cameras because of the United Serial Bus (USB) and applications are readily available on the market today, such as Intel® Create & Share™ Camera Pak and Microsoft* NetMeeting. Trends in hardware infrastructure (higher bandwidth) are evolving daily. And with a widespread base of video cameras, a higher level of combined functionality can be achieved.

Year 2000 Technology Feasibility

The feasibility of utilizing VAI technology in the year 2000 is very favorable. The segmentation, motion-detection, and object tracking, and certain types of identification technologies in VAI are in prime-time readiness status for non-mission critical applications. Let's go in greater detail to understand how these may be used in 2000 and the possibilities for the future.

Figure 1 depicts the algorithm maturity and MIPS usage of each type of VAI functionality. It is important to realize that increased algorithm maturity increases robustness to environmental changes and accuracy in VAI applications. And with increased MIP usage a higher volume platform will be required to run the application. As we discuss each technology, refer back to this chart to distinguish algorithm maturity and MIP usage.

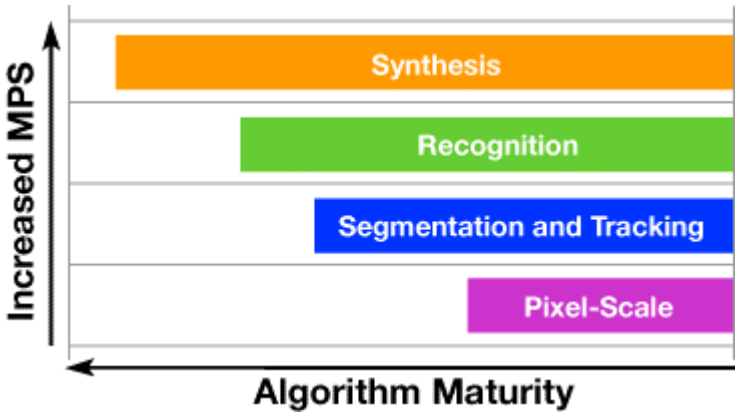

Figure 1. Year 2000 Technology Feasibility

This paper will only be focusing on the year 2000 sweet spot, that is, the basic "Pixel-Scale" and "Segmentation & Tracking" technologies depicted in the following figure.

Automatic Video Segmentation

Automatic video segmentation refers to the act of identifying portions of a video frame that belong to an object without any user intervention or special scene preparation (as in blue-screen chroma-keying used in professional productions). This task is probably the most fundamental vision task in the sense that it allows applications to interpret the video input as a collection of objects rather than a matrix of pixels. The most common use of segmentation is in applications where moving foreground objects are differentiated from stationary background objects where they are using motion as a cue. Therefore segmentation is often also called foreground/background segmentation.

From a usage model viewpoint, segmentation algorithms can be classified into two main categories. The first one being, the algorithms that require a stationary background is available before segmentation commences. The second one being the set of techniques that can detect the background in the course of segmentation. In general, the second category of algorithms is less mature and consumes greater computational resources. Consequently, applications that can afford the start-up cost of calibration generally use the first category of segmentation algorithms (see Figure 1).

Segmentation has many applications in a variety of video processing techniques including compression. Of particular interest to VAI is the use of segmentation to provide an *immersive* experience for users. In such applications, segmentation is used to extract the subject from their locale and render it into a virtual locale. This capability leads to many immersive gaming and creativity applications such as the one from Sabbatical, inc.* illustrated in the Figure 2.

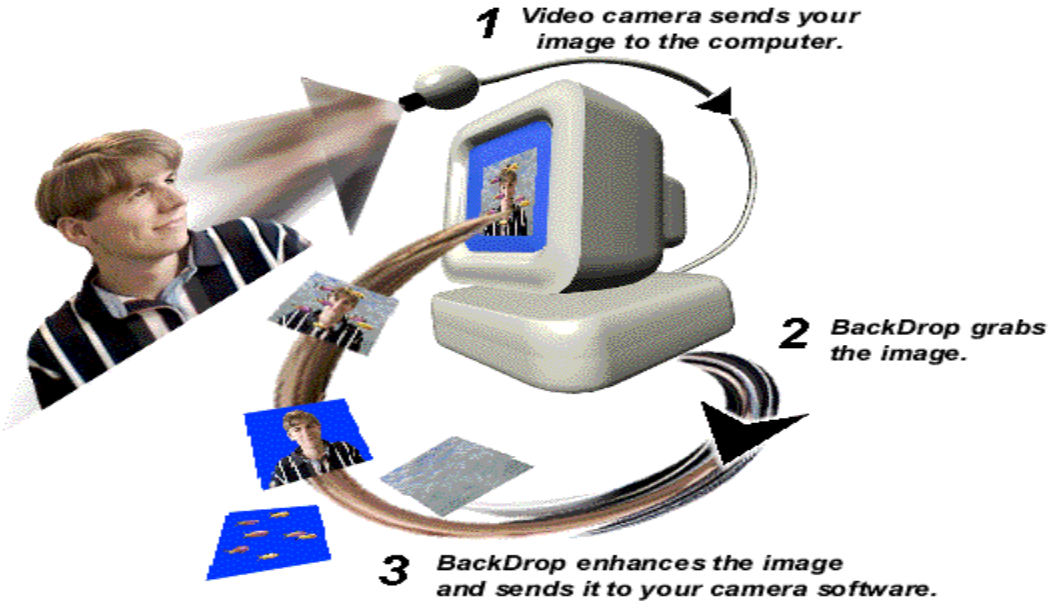

Figure 2. Sabbatical, inc. Example

Several segmentation solutions have already entered the realm of real-time operation today. Realities Fusion, Electric Planet, and Sabbatical, inc. have Software Developer Kits (SDKs) and/or applications that demonstrate the range of VAI capabilities that segmentation can enable. In general, all these implementations of segmentation can provide 20+ fps performance with 160x120 video on Intel® Pentium® II processor class machines. Off course, higher performance CPUs allow these algorithms to achieve higher frame-rate or to segment larger video frames at the same frame rate.

Motion Detection

Motion detection refers to the task of identifying temporal change in a video stream associated with object motion in specified areas of the video frame. This capability can be implemented with a range of features covering simple motion detection algorithms (providing a yes/no classification) to more capable algorithms with the ability to provide information about direction and velocity of object motion. Motion detection is frequently used to implement *hot spots* for interactive VAI applications. When combined with segmentation, motion detection can be used to provide interactive experiences in immersive (virtual) spaces. For example, a user may be immersed in a virtual room where he or she may interact with virtual objects (implemented as hot spots) using motion detection.

The simple motion detection algorithms use different forms of pixel differences; such algorithms can easily be performed in realtime with modest computing resources (Intel® Pentium® processor class PCs). The more complicated motion detection algorithms that provide speed and directional information use motion estimation or optical flow techniques, and are considerably more expensive to implement. Depending on the number of hot spots defined in the video, such algorithms may easily require the power of a Intel® Pentium® III processor class machine in order to achieve an acceptable fps (see Figure 1).

Object-Tracking and Identification

Segmentation provides VAI applications with the basic capability to address video data in terms of “objects” instead of “pixels.” An interesting class of VAI applications results from being able to detect and to react to the semantic behavior of such objects. Two such semantic tasks that are feasible today are tracking and recognition of human forms.

Tracking in the context of *human form tracking* usually refers to identifying and following interesting points of a human body such as the head, the hands, the elbow, the torso, and leg extremities. A good example of tracking technology is seen in the head-tracker capability using the Continually Adaptive Mean SHIFT (CAMSHIFT) algorithm. 2D tracking (involves tracking points in the 2D image plane of a single camera) provides the ability to drive synthetic models facilitating applications such as avatar-based communications. Such a capability can also be used to implement full or partial body gesture recognition algorithms by associating the temporal behavior of the chosen tracking points with distinct gestures. 3D tracking (where each tracked point is located in terms of a 3D space centered on the camera) usually requires stereo cameras (in order to provide depth information). While 3D tracking is more accurate and enables a wider range of application capabilities, it probably will not be deployed in volume in the year 2000, due to its dependence on stereo cameras.

Recognition is a generic task and probably most commonly used in the context of face recognition and gesture recognition (discussed above). Face recognition is an interesting and powerful technology that enables a wide variety of security (passface instead of a password) and personalization applications. A good example of commercially viable face recognition technology can be found in the FaceIt* product from Visionics. This technology can be effective on today’s PCs and promises to scale with available processing power (larger facial databases, smaller recognition times, and recognition in presence of incomplete or occluded face data) in 2000.

Break Barriers in VAI Architecture

The technology feasibility of VAI is a very real potential and although it is in it’s infancy, it has sufficient traction to create impact in the 2000 time frame. While providing exciting technologies, there is the need to be careful about creating barrier-ridden architectures. Even in its infancy, developers can take the right steps to break down wide-spread barriers. There are two types of architectures in existence today and the following descriptions will make the synergistic effort required to break down these barriers very apparent.

Monolithic Architectures

Monolithic architectures produce closed applications. This is prone to happen in cases where one developer builds the entire application stack. The result of proprietary vision algorithms is that they can only be accessed via custom APIs and custom APIs are very difficult to propagate to the enduser. This also results in vision algorithms that are tuned for specific cameras that the general population does not own or use. Another result of monolithic architectures is that only custom camera APIs can be used for controlling capture parameters that are associated with use of dated video processing infrastructures, such as Microsoft’s* Video For Window* application. Proliferation of closed applications creates barriers for the widespread deployment of VAI applications in the future.

Modular Architectures

In contrast, modular architectures create open applications. This results in applications that partition capture, vision, and display processing. The importance of partitioned tasks is that they interact via known or query-compliant interfaces, such as the video capture driver that is compliant with Microsoft’s Windows Driver Model* (WDM). Also the entire processing chain is implemented via Microsoft’s DirectShow* architecture, where vision components are filters with COM-like, query-supported interfaces.

The modular architecture advantages are that they:

- Leverage known interfaces for capture technology, where the application may work with a wider base of cameras. Custom interfaces are no longer a barrier for using a camera.
- Allow extensible use of vision technologies where minimal re-engineering to upgrade vision processing modules occurs.
- Motivate development of robust, camera-independent vision processing algorithms.
- Allow common functions to be re-used that spur faster and efficient development of a "product line." As a result, new applications do not have to be written "from scratch."

Camera Requirements
Camera interfaces as described above are important, but not the only issue. Video quality is also an important consideration. The requirements of video quality vary depending on application driven tasks. Not all cameras are created equal from a VAI perspective. In monolithic architecture cameras are created to solve the vision technology within a specific application. In other words, the camera's requirements meet task-based solutions resulting in custom interfaces. From the consumer viewpoint, a new camera would have to be purchased with every new application for optimal results. In order to provide VAI capable cameras, an understanding of vision processing is essential. This understanding will produce necessary interface hooks in the camera that will solve the "one application, one camera" problem for the future.

Video quality requirements vary depending on the task and its application. Some of the camera attributes of interest are:

Video Resolution—higher resolution provides more details for object feature detection.
- Minimum resolution for 2-feet applications is 160X120.
- Minimum resolution for 10-feet applications is 320X240.

Color—color provides essential information to many recognition tasks, e.g., face or hand detection.
- Color space is less important than color fidelity.
- Spatial sub-sampling in the sensor and interpolation must synchronize.

Field of View—must accommodate 75% of adult human body at 10 feet. Geometric distortions in lens can affect domain-knowledge-based vision tasks.

Sensor and Capture Noise—noise is a signal to many vision applications (motion detection, feature extraction).
- SNR greater than 40 db is recommended for today's technologies.
- Temporal "flickering" is an important mitigating factor.

Low Light Sensitivity—important for "home" applications. SNR greater than 40 db at 3 Lux illumination at room temperature (3K K).

Frame-rate and Driver Latency—very important for interactive applications, e.g., gesture-driven interfaces.
- 20-fps minimum; 30+ fps is recommended.
- Latency must be less than 30 milliseconds.
- Frame-rate "jitter" must be less than 5 milliseconds.

Camera Interface Requirements—programmatic interfaces must allow applications to control the camera parameters.

- Brightness
- Contrast
- Hue
- Saturation
- White Balance
- Auto Focus
- Exposure
- Frame-rate
- Color-format

Stakeholders' Call to Action

Vision Technologists & Solution Providers

Creating cameras of the future for the consumer PC can be accomplished by defining objective measures (possibly task-based) of video performance to assist camera vendors to develop "VAI-friendly" multi-function video.

Also concentrating on developing robust algorithms capable of working with different cameras and in less constrained environments, i.e. transitioning from task-based solutions to SDKs would create VAI friendly applications.

This would also include developing scalable algorithms that not only demonstrate the capabilities of tomorrow's high-end processors, but also continue to deliver enduser value across the spectrum of available computing power today. Develop algorithms written to the highest processor available that will also scale down to meet the needs of the lowest to expand the target audience.

Camera Vendors

Work with vision algorithm experts to understand and create objective video performance requirements for VAI. When advertising camera performance to application developers, publicize VAI specifications to enable the developers to build the application according to the specification. This will allow them to choose the "right" camera for the VAI application.

Abandon custom interfaces for camera drivers. Use publicized standardized interfaces that enable standardized applications. Work with interface vendors to extend interfaces (e.g. provide WDM-compliant drivers under DirectShow*).

Understand vision-processing requirements and provide all necessary interface hooks in the camera driver. One example would be to provide the ability to turn off the automatic white balancing feature.

Conclusion

Using VAI to create new ways to use visual computing has benefits for all that take the opportunity. VAI enables a class of applications that leverage the "PC + camera" to deliver new value to the enduser. Although VAI is an emerging technology area in its infancy, it has sufficient traction to create impact in the year 2000. VAI is an opportunity knocking, and tapping the full potential of VAI will require an industry-wide synergistic effort of all those who open the door to this beneficial challenge.

For More Information
The DE initiative is providing this list of vision technology solution providers to enable stakeholders and camera vendors the proof that solutions do actually exist today and are increasing quickly into 2000. This list will also provide an opportunity to glance at the next level of tool and technology solutions. Take a look at what these companies are working with and the breadth of application and usage models that are available.

Vision Technology Solution Providers:

- Electric Planet (ePlanet)
- RealityFusion
- Microsoft
- Vivid Group
- Visionics
- Sabbatical, inc.
- Intel

Designing IO Devices for a Modern PC Is Now Easy

John Hyde
Desktop Products Group
Intel Corporation

By walking you through a typical design, USB Design by Example shows that USB-based devices can be as easy to design as any other serial, parallel or ISA solution. This new book from Intel University Press contains example designs that can be tuned for specific applications and includes a companion CD-ROM with source code and schematics.

I wrote *USB Design By Example* (Intel University Press, 8/99) for everyone who wants to design an IO device for today's PC platform.

In discussions with many engineers, I discovered that they perceive USB, the Universal Serial Bus, to be a large design effort since there's so much to learn before design can even start. Many described the USB specification as a daunting study assignment. Even after reading the spec twice, many engineers found it was still not obvious how to implement a simple IO port. "I just want to pass bytes like RS232" was echoed by almost everyone. "I like USB's impressive feature set but I don't have the time and resources to invest." *USB Design By Example* dramatically changes this situation. It's the shortcut to a successful IO design project.

The Cookbook Approach to USB
USB Design By Example starts by explaining the Universal Serial Bus from an IO device's point of view. It assumes that the PC host is already operational and then it derives the expectations and responsibilities of an IO device. With this information, I define the minimum requirements of our USB-based IO device so that it can successfully attach to the PC host.

Data transfer is the next step in this cookbook. There are several approaches that can be taken, and I present the benefits of using the Human Interface Device (HID) method. The operating system contains all the device driver code for HID devices (it's well hidden but it *is* there), and by using this code you can avoid writing your own device drivers, DLLs, and other magical software.

The application programs are built step-by-step with full explanations. This approach produces simpler and more understandable solutions that are easy to modify for your specific application.

Bringing an Example to Life
By the time you reach Chapter 6, you'll have enough background to find working through a complete design example straightforward. I start with a simple "buttons and lights" example to demonstrate moving single bytes to and from the IO device.

The USB protocol "overhead" to support the IO device enumeration and then to comply with run-time requirements is implemented in less than 1 Kbyte of assembler code—including all descriptors, data, and program. While this size is large compared with the actual byte IO code, it's reusable for all the examples.

Additional examples gradually increment the complexity of the IO device so that the block mode and streaming data operation of USB can be demonstrated. These examples range from motor and lighting control, thermometers, modems, and audio up to a full TV-tuner dongle. Low- to medium-speed devices are implemented as HID, while higher data-rate devices are implemented using other USB class types.

After reading the book, you'll know that a simple USB-based IO device is as easy to implement as a serial, parallel, or ISA-based solution. You'll also find enormous headroom in the USB solution: USB supports block transfers, streaming data, and multiple "channels" that in turn support command/data, configuration/run time, and other constructs. These constructs make your IO device easier to design and easier to use. And your IO device automatically inherits all the benefits of USB.

If you've been thinking about USB but haven't had the time or resources to invest in a large engineering project, you can select one of the working examples from this book and tune it to meet your needs. Your custom IO device will be complete before you know it.

About the Author

John Hyde is a 22-year Intel veteran who has been responsible for providing technical training and technical documentation for a wide range of Intel's products. You have probably already read some of his, or his groups, work! He was technical marketing manager for the Multibus* II product line and for the Intel® Pentium® Pro microprocessor.

John realized the need for "USB Design By Example" while he was Intel's editor for the PC98 System Design Guide – this joint Intel/Microsoft document started the outlaw of the ISA bus and the serial and parallel ports since they have such a detrimental effect on the PC's system performance. USB was the recommended solution but the PC industry proclaimed that this was "just too difficult to use".

John obtained his Bachelors Degree in Electronics from Southampton University in England. His professors wanted him to complete their Masters Program but John, not wanting to be Dr. Hyde, joined Intel as a field applications engineer instead. His first project was the support of a 5 volt-only version of the 8080A microprocessor, later called the 8085, and he has been learning ever since!

For More Information

All source code and schematics are included on a companion CDROM and are supported by a companion Web site. The PC host software is written in Visual Basic, and the IO device firmware is written in MCS-51 assembler code. The CDROM includes MCS-51 development tools you can use to modify each example to suit your requirements.

Intel® FlexATX Form Factor Helps Reshape the PC

Frank Soqui
Advanced PC Platform Development Manager
OEM Platform Solutions Division
Intel Corporation

Are you ready for the "Barbie PC" and the "Hot Wheels* PC"? Intel's new FlexATX motherboard form factor ushers in a new era of innovative PC design and novel product marketing opportunities for OEMs. Preview the cool new products based on FlexATX boards at the Connected Home Demo at IDF.*

Are you ready for the “Barbie PC”? How about the “Hot Wheels PC”? Intel’s new FlexATX motherboard form factor ushers in a new era of innovative PC design, in addition to some imaginative product marketing opportunities for OEMs. You can see a variety of new products based on FlexATX boards at the Connected.Home demo at IDF and featured in the “Designing Next Generation Performance Desktop Systems” track.

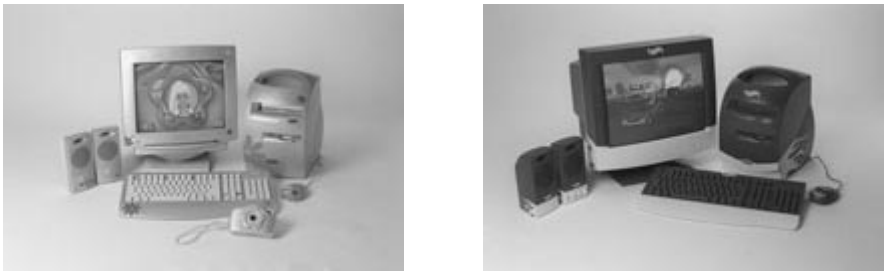

Barbie PC and Hot Wheels PC

For customers who have grown tired of “beige box” PCs, Intel’s FlexATX motherboard form factor represents a big (or should we say small?) step in a new direction. Measuring just 7.5 x 9.0 inches, FlexATX desktop boards are the ideal foundation for creating smaller-sized PCs, consumer-friendly closed-box PCs, and new industrial designs.

FlexATX has already helped generate a host of novel PC designs, including AST’s Easy PC* and the new Barbie PC and Hot Wheels PC from Patriot Computers. You can see the latest PC adaptations at the Connected Home Demo at Fall ’99 Intel Developer Forum. These designs are sure to spark your imagination. Technical details will be presented in the session “Platforms: The Shape Of Things to Come,” which will give you the information to enable you to base your industrial designs on this exciting new form factor.

Patriot’s new products test the waters for “lifestyle” PCs, with designs targeted at specific demographic groups often overlooked in PC marketing. Together with AST’s new offering, these computers provide the latest examples of new design directions made possible by Intel’s new motherboard specification. The PCs are designed with simple, highly integrated Intel® Architecture components designed to enhance ease-of-use, while keeping development costs and product price points as low as possible.

Good Things Come in Small Packages
Intel’s FlexATX motherboard products support Intel® processors in the FCPGA form factor and feature the Intel® 810 and family of chipsets with integrated 3D graphics. While some are specifically designed for Intel® Celeron™ processors, other FlexATX boards scale to Intel® Pentium® III processors.

Intel is developing a wide variety of configurations for FlexATX desktop boards. Some of the desktop boards have been specifically designed for closed-box designs that rely on external expansion and feature no PCI slots. End-users expand these systems through as many as five USB ports at the front and back panels for simple plug-and-play expandability. Some FlexATX desktop boards have eliminated legacy connectors including PS/2, MIDI/game port, serial and parallel ports.

Defined as an addendum to the microATX Specification, the FlexATX form factor offers some very attractive advantages for developers:

Compatibility with ATX and microATX chassis provides a high level of design flexibility and support for multiple chassis in one desktop board.

With greater integration of features such as modem and HPNA-based home networking, PCI expansion cards are unnecessary on some versions of Intel's FlexATX desktop boards. The bottom line for developers is that systems with eye-catching styling and breakthrough industrial design can be made ready to go with a minimum of technical design effort.

Innovative technologies, such as the Digital Video Interface (DVI) specification for flat-panel displays, support the transition to new industry standards.

The use of readily available components, such as standardized power supplies, helps keep development costs to a minimum.

It's Time to Get Creative
With building blocks from Intel like the new FlexATX form factor, the Easy PC Hardware Implementation Guide, and Intel's support for chassis manufacturers to create new industrial designs, Intel is working to help developers large and small implement new, innovative designs that support the latest technologies.

By combining small size, ATX-based compatibility, high integration, standardized components, and relatively low cost, the FlexATX form factor sets the stage for a variety of exciting new PC-based products. See what's new at IDF, and then prepare to turn your imagination loose.

About the Author
In his role as APPD marketing manager, Frank J. Soqui focuses on development and implementation of marketing strategy for desktop form factor enabling, as well as system and motherboard products in support of those enabling efforts.

A seventeen-year veteran of Intel, Frank was previously product marketing manager for desktop motherboard products for OEM Platform Solutions Division in DPG (Desktop Products Group). He has also served as product marketing engineer for several communications products including modem components, branded modem products, and software-based modem solutions. Prior work also included silicon and software engineering at Intel. He received his Bachelor of Science in Electronic Engineering Technology from DeVry Institute of Technology.

For More Information

- Check out Intel's Easy PC Initiative Web site.
- Visit the Platform Development Support Web site for specifications and the latest information on FlexATX.
- For all of Intel's desktop board offerings, be sure to visit the developer and channel Web sites.

Technology News Bytes

August 23

Intel Announces Latest Pentium® III Xeon™ Processors

Offering Higher Levels of Performance, Headroom and Scalability; Profusion® Chipset for 8-Way Enterprise Servers also Shipping.

Intel Corporation announced the latest versions of its Pentium® III Xeon™ processor family for server and workstation platforms. The three new 550 MHz processors incorporate 512 KB, 1 MB and 2 MB of Level 2 (L2) cache, which support four-way and greater server and workstation configurations. Intel also announced that it is shipping its Profusion chipset, enabling eight-way Pentium III Xeon processor based servers.

New Intel® PC Camera Directly Connects Camcorders, VCRs to PCs

Intel introduced the Intel® PC Camera Pro Pack, an affordable and easy-to-use package that includes video phone, video e-mail, and auto snapshot capabilities, and is the first PC camera with a built-in video capture plug to bring live or recorded video directly into PCs.

Intel Science Talent Search Awards Increase to $1.2 Million a Year;

Semifinalists and Schools to Benefit from Awards for the First Time in the 59-year History of the Program.

Intel Corporation and the Intel Foundation announced a grant to Science Service, the administrator of the Intel Science Talent Search (Intel STS), to increase the yearly awards for the program from $330,000 to $1.2 million a year, beginning this year. Total awards for the Intel STS have increased by $1 million since 1998 when Intel became the sponsor of the prestigious pre-college science competition, formerly sponsored by Westinghouse.

August 16

Intel Makes Shared Internet Access Simple and Affordable for Small Businesses;

Intel® InBusiness™ Internet Station 56K Features Built-in Modem and Hub; Small Business Can "Connect" for Less Than $300.

Intel introduced the Intel® InBusiness™ Internet Station 56K, a network appliance designed to help small businesses connect to the Internet simply and affordably. It allows several employees to access the Internet from a single account. About the size of a paperback novel, the Intel InBusiness Internet Station 56K combines a 56 Kbps (kilobits per second) V.90 analog modem with a four-port 10 Mbps network hub for easy sharing of office resources and Internet access.

August 10

Intel Completes Merger with Level One Communications

Intel announced the completion of its merger with Level One Communications, Incorporated, following the approval of Level One shareholders. As a result of the approval, each Level One share is converted into .86 shares of Intel stock. Level One will operate as a wholly owned subsidiary of Intel.

Intel's Maloney Delivers Linux World Keynote;

Highlights Intel's IA-64 Architecture as the New Engine for E-Business; Discloses New Developer Initiatives.

At Linux World, Intel Senior Vice President Sean Maloney delivered the keynote presentation, highlighting IA-64, Intel's forthcoming 64-bit server and workstation architecture, as the new engine for e-business in next-generation computing. Maloney also disclosed several new initiatives for developers of operating systems targeted to IA-64, including Linux. He was joined on stage by Intel Chairman Andrew S. Grove, for a demonstration of the Linux operating system running on an IA-64 software development environment.

IDT And Intel Sign Cross-Licensing Agreement
IDT, Inc. (Nasdaq: IDTI) and Intel announced the two companies have entered into a cross-license agreement that enables each company to utilize the intellectual property (IP) covered by the other's patents.

August 9
Hollywood's Digital Domain Moving to Intel® Architecture for Content Design and Creation
Digital Domain, one of Hollywood's leading digital visual effects studios and effects creators for such movies as "Titanic," "What Dreams May Come," and "The Fifth Element," today announced its plan to migrate to Intel® Architecture based workstations. The move will be based on the dual Pentium® III and Pentium® III Xeon™ processor platforms. The new systems will be integrated with existing Intel based workstations that Digital Domain has been using to produce special visual effects for feature films, commercials and new media projects. As a further step in the migration, Digital Domain intends to serve as a test site for Intel's IA-64 family of processors.

Intel Announces Systems Management Support for Linux*;
Intel® LANDesk® Management Suite 6.3 Extends to Provide Remote Access and Inventory for Linux-based Systems.

Intel announced it has extended the Intel LANDesk Management Suite to provide support for Linux-based systems. A new software module for Intel LANDesk Management Suite 6.3, available today for download from the Internet, enables IT managers to inventory the hardware and software assets on Linux-based laptops, PCs and servers and take remote control of those systems when diagnosing or repairing problems.

Intel Broadens Networking Line with Easy-to-Use, Manageable Switch;
Express Hub and Switch Price Cuts Make Intel's Comprehensive Line More Affordable.

Intel announced the Intel® Express 460T Standalone Switch, completing its product line rollout to offer customers a full range of LAN networking solutions for each segment in the standalone switch category. The Express 460T switches combine intuitive control and configuration with built-in manageability and ease-of-use features. Intel also announced an 8 to 22 percent price reduction across its Express switch and hub product line.

August 2
Intel Corporation Announces 3 Volt Intel® Strataflash™ Memory;
NOR Flash Offers Faster Access to Improve the Functionality of Connected Appliances.

Intel, the leader in flash memory, announced today its 3 Volt Intel® StrataFlash™ memory with triple the read performance over the previous version. Using advanced 0.25-micron lithography, Intel's memory enables both code execution and data storage on a single high density 128 Mb chip.

Intel Ships Fastest Intel® Pentium® III, Intel® Celeron™ Processors;
New Products Arrive for Back-to-School PC Buying Season.

Intel introduced two new processors, the Intel® Pentium® III processor 600 MHz for powerful Internet and mainstream computing, and the Intel® Celeron™ processor 500 MHz for sub-$1,000 PCs. The new Intel Celeron processor 500 MHz is Intel's fastest processor for value PCs.

August 1
Intel and Pacific Century Group Announce Agreements
to Enable Broadband Internet Deployment in Asia; Intel to Invest US$50 Million in Pacific Century CyberWorks; Intel to Supply Comprehensive Set Top Solutions.

Intel and Hong Kong-based Pacific Century Group (PCG) announced agreements to enable broadband Internet deployment in Asia. The announcement comprises three related elements, subject to the completion of PCG's acquisition of a controlling interest in Tricom Holdings Limited, which will subsequently be renamed Pacific Century CyberWorks Limited.

Industry Events

Industry Events for September

XDC 1999
September 8
Wembley Conference Center
London, U.K.
XDC is a special event for top Intel® Architecture Developers. In 2000, PCs will be more powerful than ever, with unprecedented processor power, stunning 3D graphics and fast Internet connectivity. Intel is committed to help software developers take full advantage of all the hottest new hardware technology to create innovative and exciting software. Intel's annual XDC events provide both consumer and business software developers with the opportunity to learn what is required to develop outstanding applications for current and future IA-based PC platforms. Intel and industry experts demonstrate how to apply the latest technology and programming techniques to create the cutting-edge, market-leading software for the Internet as well as for standalone applications. Attendees will learn how to optimize applications for Intel® Pentium® III processors and beyond; get the best graphics and processor performance; make the right tools and engines work for their applications.

System Builder Summit
September 9
Boca Raton Resort and Club
Boca Raton, FL
This is an exclusive forum that brings together the powerhouses in the white-box market with leading technology vendors. In 1999, white-boxes will account for 42 percent of the sales of new PCs. The System Builder Summit hosts the top 1 percent of the leading system builders who represent significant purchasing power. System builders, who attend as guests of the event, discuss critical channel issues with each other and with senior executives from manufacturing, software, distribution and service organizations.

Boon Chye Ooi, Vice President, Intel Desktop Products Group, and General Manager, OEM Platform Solutions Division, keynotes at SBS.

NetWorld+Interop 99
September 13-17
Georgia World Congress Center
Atlanta, GA
Mark Christensen, Vice President and General Manager, Intel Network Communications Group, keynotes Tuesday, September 14 from 5:00 to 6:00 p.m.

Integrating Applications 1999
(formerly Middleware Choices Europe)
September 21-24
Millennium Centre
London, U.K.
As e-commerce becomes a reality, there is a fundamental transformation in the way business is conducted. Businesses are rapidly moving to global commerce—electronically linking customers and suppliers around the world. Internet applications are no longer confined to the "edge" of business applications. Today, every part of the business and associated applications has either already been touched by the Internet or is about to be. Giga Information Group's Integrating Applications conference is designed to help make sense of the rapidly evolving EAI market, defining the options and offering guidance through critical technical and managerial challenges. Intel Vice President and General Manager, Communications Products Group John Miner addresses the conference on September 23.

Embedded Systems Conference
September 27-30
San Jose Convention Center
San Jose, CA
Intel's Computing Enhancement Group and its Applied Computing products and technologies will be represented in a booth. Vice President and General Manager Ron Smith will be a featured speaker.

1999 NGIO Forum's Developers' Conference
September 28-30
Newport Beach, CA
September 28-30
Newport Beach, CA
The conference focuses on Next Generation I/O (NGIO) and offers two and a half days of in-depth sessions. Particular attention will be given to design implementation. Intel is a steering committee member of the NGIO Forum and will have several experts on hand to present the latest information for developers. For information on NGIO servers visit Intel's Next Generation of I/O Servers web site.

Industry Events for October

Embedded Internet Workshop
October 1
San Jose, CA
Tom Franz, Vice President and General Manager, Intel Embedded Microcomputer Division, is the keynote speaker who opens the conference.

Fall Internet World '99
October 4-8
New York, NY
Whether you're a webmaster evaluating E-commerce software, an ISP rethinking business strategy, a marketing executive looking for proven online opportunities, or an IT professional who needs to secure an enterprise network, Internet World has something for you. Billed as “The World's Largest Event for E-Business and Internet Technology,” Fall Internet World 99 rivals the best stand-alone, vertical conferences in the industry for both content and emerging leadership. From advanced technology workshops to introductory sessions, attendees can choose from a range of programs covering every aspect of conducting business on the Internet.

Intel is the corporate sponsor of the conference. Intel President and CEO Craig Barrett gives the opening keynote 9:00-9:45 a.m. Wednesday, October 6.

Gartner Group U.S. Symposium/ITxpo 99
October 12
Lake Buena Vista, FL
Symposium/ITxpo 99 will address the entire spectrum of IT with the multiple needs of business in mind -- now and in the future. If you plan, manage, strategize, guide, purchase, market, direct or use IT, Symposium/ITxpo is a crucial strategic tool.

The Symposium will conduct a Mastermind Interview with Intel President and CEO Craig Barrett. In this interview Gartner will ask Dr. Barrett about Intel's view of global connected computing, advanced networking, distributed vs. central computing, and the future of the microprocessor. Gartner analysts will question him on Intel's strength, strategies, partnerships, and competition.

United States Telephone Association 102nd Annual Convention & Exhibition
October 17-19
Hyatt Regency Hotel
San Francisco, CA
The USTA Annual Convention brings together senior executives from all segments of the telecommunications industry. Telecom executives attend the annual event to network with other industry professionals, participate in high-level discussions about critical industry issues and to gain valuable insights that will help them grow their businesses. Over 1,000 participants are expected to attend the Convention, representing every area of the communications industry. This year’s theme is “Communications Convergence: Focus on Customer Solutions.”

The opening keynote speakers will be Intel co-founder and chairman Andy Grove and FCC Chairman William Kennard.

Internet Commerce Expo (ICe)
October 18-21
Moscone Convention Center
San Francisco, CA
Internet Commerce Expo is the leading eBusiness trade show and conference focusing on Internet Commerce, intranets, extranets and Web-based business solutions. ICe is specifically targeted at business users and IT professionals that are buying Web-based technologies and applications to advance productivity and profitability gains in the eBusiness enterprise.

Intel will be exhibiting at ICe. Paul Otellini, executive vice president and general manager, Intel Architecture Business Group, will present a keynote address on October 20, 9:30 – 10:30 a.m.

DV Expo 99
October 19-22
Long Beach Convention Center
Long Beach, CA
DV Expo is the place to meet peers and experts, learn time-saving techniques, and get the low-down on the latest tools and technologies. The Digital VideoConference & Expo offers four intense days of in-depth training on digital video technology and business. It puts attendees on the fast track for learning the latest tips and tricks and evaluating the best product solutions for digital video projects. Topics include production, editing, motion graphics, 3D, visual effects, corporate video, compression, DVD and Web, and more. Intel is an exhibitor.

ISPCON Fall '99
October 26-28
San Jose Convention Center
San Jose, CA
ISPCON Fall '99 will gather thousands of Internet Service Providers, CLECs, ICSPs, ITSPs, IXCs, Cable Operators, Next Gen Telcos, Telecom Resellers, Satellite/Wireless Companies, and Web Hosting companies for an intensely educational event focused on the business of providing Internet access.

Intel is an exhibitor and corporate sponsor. Communications Products Group vice president and general manager John Miner will be speaking during the conference.

Industry Events for November

Internet World Asia@Hong Kong '99
November 8
Hong Kong Convention & Exhibition Centre
Hong Kong
Since 1993 Internet World has consistently captured the state-of-the-art in Internet technology. That's why forward-thinking business and IT professionals converge on Internet World to learn about advances in Internet technologies and products as well as the companies behind them. Internet World is an important tradeshow and conference for professionals whose careers and businesses now depend on the Internet for their future success.

Intel is an exhibitor. In addition, Executive Vice President and General Manager, Intel Architecture Business Group Paul Otellini keynotes.

Oracle Open World
November 14-18
Los Angeles, CA
The largest technology conference of its kind, Oracle OpenWorld brings together a unique group of Oracle customers, the industry's top analysts and technicians, and Oracle executives and senior product experts. It provides the high-quality information businesses need to meet real-world technical challenges. The year's best chance to network with peers and decision makers, it's an opportunity to get answers from the people who build Oracle products and the people who use them to run their businesses. Oracle OpenWorld is designed and developed by Oracle Corporation in cooperation with the International Oracle Users Group - Americas (IOUG-A).

Intel President and CEO Craig Barrett keynotes on November 16.

—End of Platform Solutions Issue 23—